

Resource-efficient and robust quantum federated learning via loss aggregation

Cheng-En Tsai^a, and Hao-Chung Cheng^{a,b,c,d,e,*}, ORCID

Abstract

Federated learning across quantum devices offers a pathway to scalable and privacy-preserving quantum machine learning. However, existing frameworks exchange large gradients or model parameters, creating severe communication and execution overheads that scale with system size and often destabilize training on heterogeneous client data. We introduce Loss Aggregation Quantum Federated Learning (LossAgg-QFL), a framework that transmits only scalar loss values from each client to a central server. Using Simultaneous Perturbation Stochastic Approximation (SPSA), the server infers global parameter updates through constant-cost evaluations, eliminating dependence on qubit number, circuit depth, and parameter dimension. This design reduces per-round communication and circuit executions by up to 10^6 times and improves runtime scaling from $\mathcal{O}(n^3)$ to $\mathcal{O}(n)$ with respect to qubit count. A multi-directional SPSA strategy further stabilizes optimization, lowering convergence variance from 3.6% to 0.8% across stochastic, gradient-free updates. Evaluations on the MNIST and Fashion-MNIST benchmarks show that LossAgg-QFL achieves up to +59% and +33.8% improvements in accuracy under label- and quantity-skew conditions, and also outperforms baseline methods under balanced data. By aligning learning dynamics with quantum-hardware constraints and federated scalability, LossAgg-QFL demonstrates that loss-based aggregation can elevate quantum federated learning from primarily simulation-driven studies to a more practical and resource-efficient paradigm for distributed quantum intelligence.

I. INTRODUCTION

Quantum Federated Learning (QFL) [1]–[4] combines quantum computing [5], [6] with the decentralized architecture of Federated Learning (FL) [7], [8], offering a promising path toward scalable, privacy-preserving distributed quantum intelligence. As hardware roadmaps progress toward modular, distributed, and multi-chip quantum processors [9], [10], QFL naturally aligns with these interconnected resources. However, practical deployment is severely hindered by several critical bottlenecks.

A central barrier is *communication overhead* [11], [12]. Conventional frameworks exchange large parameter vectors or gradient sets with costs that grow with qubit count and model size, creating prohibitive bandwidth usage and complicating privacy-preserving protocols [13], [14]. Another fundamental challenge is the *non-IID quagmire* [15]–[17], where heterogeneous, imbalanced, or label-skewed client data cause high-dimensional local gradient updates to magnify local discrepancies, resulting in training divergence and reduced accuracy [18], [19].

These obstacles are further exacerbated by a gap between simulation and hardware feasibility. Most QFL studies evaluate with classical simulators using exact-gradient optimizers such as Adam [20]. On real quantum hardware, however, back-propagation [21] is infeasible due to wavefunction collapse [22], [23]. While the parameter-shift rule [24], [25] provides accurate gradient estimates, its circuit executions scale extensively with model size and depth, making repeated evaluations under device noise unstable and resource-prohibitive [26]. Taken together, these limitations highlight the need for a new framework that is both *resource-efficient* in hardware cost and *robust* to heterogeneous data.

To this end, we propose *Loss Aggregation Quantum Federated Learning (LossAgg-QFL)*, a resource-efficient and robust framework designed for practical quantum processors. Instead of transmitting large parameter vectors, clients report only scalar losses, and the central server estimates gradients using a multi-directional variant of Simultaneous Perturbation Stochastic Approximation (SPSA) [27], [28]. This design reduces communication and execution costs by orders of magnitude, while enhancing robustness against non-IID data and optimization stochasticity. Our core contributions are:

^aDept. of Electrical Engineering, National Taiwan University, Taipei, Taiwan; ^bDept. of Mathematics, NTU; ^cCenter for Quantum Science and Engineering, NTU; ^dHon Hai (Foxconn) Quantum Computing Center, New Taipei City, Taiwan; ^ePhysics Division, National Center for Theoretical Sciences, Taipei, Taiwan. *Corresponding author. A full standalone version of this manuscript is available for reference.

- **Resource efficiency in quantum federated learning:** LossAgg-QFL reduces communication (uplink data) and executions (circuit evaluations multiplied by shots) to a constant overhead $\mathcal{O}(1)$ per client per training round, independent of qubit count, circuit depth, or model architecture. It lowers device runtime scaling from $\mathcal{O}(n^3)$ to $\mathcal{O}(n)$ with respect to qubit number n [29], [30], yielding up to a $10^6\times$ total reduction.
- **Robustness to data heterogeneity and stochastic updates:** By aggregating scalar losses rather than high-dimensional gradients, LossAgg-QFL mitigates bias from heterogeneous client updates under non-IID distributions, achieving up to +59% accuracy gains over FedAvg [8] and FedAdam [18] in label-skew scenarios. Furthermore, its multi-directional SPSA strategy reduces gradient-free estimation variance from 3.6% to 0.8%, ensuring stable convergence.

II. PROPOSED FRAMEWORK

We introduce Loss Aggregation Quantum Federated Learning (LossAgg-QFL)¹, a resource-efficient and robust framework designed to bypass high-dimensional communication overhead and the non-IID data heterogeneity quagmire. Conventional QFL frameworks exchange large gradient or parameter tensors between clients and the central server, an approach that is both communication-intensive and incompatible with quantum hardware where back-propagation is infeasible due to wavefunction collapse [22]. LossAgg-QFL reframes this paradigm by transmitting only *scalar loss values* from each client to the central server, decoupling learning efficiency from parameter dimensionality.

The protocol operates via a decoupled, loss-centric communication loop across four core phases: (1) The central server samples and broadcasts N_{SPSA} random perturbation trajectories $\{\Delta_t^{(n)}\}$ to all participating clients; (2) Local clients evaluate perturbed circuit losses $L_i^{(n)\pm} = L(x_t \pm c_t \Delta_t^{(n)})$ strictly on their private datasets D_i , avoiding local gradient or parameter updates; (3) The server aggregates these scalar responses into global averages $\bar{L}_n^\pm$ weighted by dataset size; and (4) The server reconstructs the global descent trajectory from the aggregated scalars, executing global parameter updates via multi-directional stochastic approximation.

Each training round computes global weighted loss averages $\bar{L}_n^\pm = (\sum_i D_i)^{-1} \sum_i D_i L_i^{(n)\pm}$, which act as unbiased estimators of the total loss across the system. Rather than relying on parameter-shift rules scaling as $\mathcal{O}(p)$, LossAgg-QFL reconstructs the global gradient using a multi-directional stochastic approximation strategy. For each perturbation direction n , the central server approximates the global gradient vector as: $g_t^{(n)} \approx \frac{\bar{L}_n^+ - \bar{L}_n^-}{2c_t} \Delta_t^{(n)-1}$, where $\Delta_t^{(n)}$ is a random perturbation vector drawn from a Rademacher distribution (± 1), enabling the simultaneous approximation of all high-dimensional parameter trajectories using only two loss evaluations per sampled direction.

To mitigate quantum sampling fluctuations, transient device noise, and non-IID variance, we average these reconstructed estimates over N_{SPSA} random trajectories. The central server integrates these updates using two distinct aggregation options: **LossAgg (Averaging)**, which averages updates across directions $x_{t+1} = \frac{1}{N_{\text{SPSA}}} \sum_n x_t^{(n)}$ following FedAvg principles for low-overhead parameter estimation; or **LossAggAdam (Adaptive)**, which incorporates momentum tensors $m_t = \beta_1 m_{t-1} + (1 - \beta_1)(x_{t+1} - x_t)$ and variance tracks $v_t = \beta_2 v_{t-1} + (1 - \beta_2)(x_{t+1} - x_t)^2$ with bias-corrections to dynamically update global parameters via $x_{t+1} = x_t + \eta \hat{m}_t / (\sqrt{\hat{v}_t} + \tau)$. By communicating only scalar values, this design limits structural information exposure, natively mitigating vulnerability to gradient inversion attacks [31] without altering the core communication primitive.

The framework was evaluated using Qiskit [32] across Iris, MNIST [33], and Fashion-MNIST (FMNIST) [34] benchmarks, staging 10-client and 8-client setups with a mini-batch size of 16. To simulate real-world heterogeneity, three canonical partitioning strategies were tested: (i) *Label distribution skew*, restricting clients to $\#C \in \{1, 2\}$ classes; (ii) *Quantity skew*, generated via a Dirichlet distribution ($\alpha = 0.5$); and (iii) *IID* uniform distribution. Performance evaluation metrics consist of uplink payload volume, total circuit executions (circuits $\times$ shots), wall-clock device occupation runtime, and top-1 classification accuracy across 5 independent runs.

III. MAIN BENCHMARKS

A. Communication, Execution, and Runtime Complexity

For a QNN with n qubits, L layers, and c trainable parameters per qubit per layer ($p = c \cdot n \cdot L$), practical ansätze require parameterized two-qubit gates scaling quadratically with width ($c = \mathcal{O}(n)$) [35], while depth L

¹All code and implementation protocols are open-source at <https://github.com/xNTsai/LossAgg-QFL>.

scales linearly with n [36]. Consequently, conventional parameter-shift (PS) frameworks incur a per-round uplink communication payload of $C_{\text{PS}} = \mathcal{O}(p) = \mathcal{O}(n^2 L) \rightarrow \mathcal{O}(n^3)$. LossAgg transmits only two scalar loss values per perturbation trajectory direction. Given N_{SPSA} trajectories, the per-round uplink overhead remains completely flat at $C_{\text{LossAgg}} = 2N_{\text{SPSA}} = \mathcal{O}(1)$, fully independent of circuit width or layer depth, yielding a $10^3 \times$ reduction at $(n, L) = (64, 50)$. This structural separation extends to physical quantum hardware execution counts (N) and device occupation runtimes (T). Standard parameter-shift protocols require two independent circuit evaluations per parameter per training update direction, yielding $N_{\text{PS}} = 2pRS = \mathcal{O}(n^2 L)$. LossAgg requires only two evaluations per perturbation direction: $N_{\text{LossAgg}} = 2N_{\text{SPSA}}S = \mathcal{O}(1)$. At moderate circuit complexities $(n, L) = (32, 40)$, PS requires over 10^7 total circuit executions per round, whereas LossAgg remains pinned at a constant baseline, yielding a $10^6 \times$ resource reduction. Accounting for layer-wise hardware scheduling overhead where gate depth D scales with n , the hardware execution runtime shifts from $T_{\text{PS}} = \mathcal{O}(n^3 L^2)$ down to $T_{\text{LossAgg}} = \mathcal{O}(nL)$, cutting active hardware occupation time by more than $10^5 \times$ at $(n, L) = (32, 40)$.

B. Variance Suppression and Heterogeneous Data Resilience

While single-direction gradient-free stochastic approximation minimizes evaluation complexity, it can introduce pronounced stochastic variance under quantum sampling noise, thereby destabilizing optimization trajectories. LossAgg-QFL suppresses these fluctuations by averaging gradient estimates across N_{SPSA} orthogonal trajectories per round, smooths the global descent vector, and enables higher learning rates without expanding communication volume beyond its strict $\mathcal{O}(1)$ boundary. Empirically, scaling N_{SPSA} from 10 to 40 smooths the optimization trajectory, reducing convergence variance from 3.6% to 0.8% while elevating top-1 mean accuracy from 48.9% to 83.5%.

To evaluate resilience against non-IID heterogeneity, extensive empirical benchmarks were conducted across Iris, 8-class, and 10-class variants of MNIST and Fashion-MNIST under diverse partitioning topologies (including label skews, quantity skews, and IID controls). Table I presents representative, highly constrained non-IID configurations from these evaluations to isolate optimization stability under extreme statistical distribution shifts.

TABLE I
TOP-1 CLASSIFICATION ACCURACY (%) UNDER HETEROGENEOUS DATA CONFIGURATIONS

Dataset	Non-IID Partitioning Strategy	FedAvg	FedAdam	LossAgg	LossAggAdam
10-MNIST	Label Skew ($\#C = 1$)	18.6 ± 3.8	30.9 ± 2.7	68.8 ± 0.9	70.1 ± 3.1
	Quantity Skew ($q \sim \text{Dir}(0.5)$)	55.4 ± 2.5	61.4 ± 0.5	67.5 ± 0.6	67.7 ± 3.0
10-FMNIST	Label Skew ($\#C = 1$)	20.5 ± 6.0	31.4 ± 4.9	63.3 ± 1.7	62.0 ± 1.7
	Quantity Skew ($q \sim \text{Dir}(0.5)$)	58.0 ± 3.2	59.0 ± 2.1	63.0 ± 1.5	61.3 ± 1.8

As summarized in Table I, high-dimensional gradient-exchange baselines (FedAvg and FedAdam) exhibit severe performance degradation and early convergence plateaus due to local update updates amplifying client discrepancies. In contrast, by aggregating low-dimensional scalar loss values instead of large parameter tensors, LossAgg-QFL filters out client-specific parameter bias. In extreme label-skew settings, LossAgg outpaces FedAvg by up to **+59%** and FedAdam by **+48.1%**. Under quantity-skew configurations, it yields up to a **+33.8%** accuracy gain, establishing uniform convergence across heterogeneous networks.

IV. CONCLUSION

LossAgg-QFL introduces a paradigm shift in distributed quantum machine learning by substituting high-dimensional parameter or gradient tracking with a scalar loss aggregation primitive. Driven by a multi-directional stochastic approximation engine, the framework reduces per-round uplink communication bandwidth, discrete circuit executions, and physical hardware runtime complexities by up to six orders of magnitude. Crucially, by isolating and neutralizing local update updates, it resolves the non-IID heterogeneity quagmire, yielding substantial classification accuracy gains under severe label and quantity skews. This combination of constant-cost hardware scaling and data resilience transforms quantum federated learning from a simulation-bounded theory into a highly practical protocol for real-world distributed quantum processors.

REFERENCES

- [1] Rui Huang, Xiaoqing Tan, and Qingshan Xu. Quantum federated learning with decentralized data. *IEEE Journal of Selected Topics in Quantum Electronics*, 28(4: Mach. Learn. in Photon. Commun. and Meas. Syst.):1–10, 2022.
- [2] Qi Xia and Qun Li. Quantumfed: A federated learning framework for collaborative quantum training. In *2021 IEEE Global Communications Conference (GLOBECOM)*, pages 1–6, 2021.
- [3] Chao Ren, Rudai Yan, Huihui Zhu, Han Yu, Minrui Xu, Yuan Shen, Yan Xu, Ming Xiao, Zhao Yang Dong, Mikael Skoglund, Dusit Niyato, and Leong Chuan Kwek. Toward quantum federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, 36(9):15580–15600, 2025.
- [4] Mahdi Chehimi and Walid Saad. Quantum federated learning with quantum data. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8617–8621, 2022.
- [5] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018.
- [6] Shubham, Prachi Sajwan, and N. Jayapandian. Challenges and opportunities: Quantum computing in machine learning. In *2019 Third International conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud) (I-SMAC)*, pages 598–602, 2019.
- [7] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020.
- [8] H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data, 2023.
- [9] Muhammad AbuGhanem. Ibm quantum computers: evolution, performance, and future directions. *The Journal of Supercomputing*, 81(5), April 2025.
- [10] D. Main, P. Drmota, D. P. Nadlinger, E. M. Ainley, A. Agrawal, B. C. Nichol, R. Srinivas, G. Araneda, and D. M. Lucas. Distributed quantum computing across an optical network link. *Nature*, 638(8050):383–388, February 2025.
- [11] Osama Shahid, Seyedamin Pouriyeh, Reza M. Parizi, Quan Z. Sheng, Gautam Srivastava, and Liang Zhao. Communication efficiency in federated learning: Achievements and challenges, 2021.
- [12] Harashta Tatimma Larasati, Muhammad Firdaus, and Howon Kim. Quantum federated learning: Remarks and challenges. In *2022 IEEE 9th International Conference on Cyber Security and Cloud Computing (CSCloud)/2022 IEEE 8th International Conference on Edge Computing and Scalable Cloud (EdgeCom)*, pages 1–5, 2022.
- [13] Changhao Li, Niraj Kumar, Zhixin Song, Shouvanik Chakrabarti, and Marco Pistoia. Privacy-preserving quantum federated learning via gradient hiding. *Quantum Science and Technology*, 9(3):035028, May 2024.
- [14] Zhonghui Li, Kaiping Xue, Jian Li, Lutong Chen, Ruidong Li, Zhaoying Wang, Nenghai Yu, David S. L. Wei, Qibin Sun, and Jun Lu. Entanglement-assisted quantum networks: Mechanics, enabling technologies, challenges, and research directions. *IEEE Communications Surveys & Tutorials*, 25(4):2133–2189, 2023.
- [15] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated learning with non-iid data, 2018.
- [16] Kevin Hsieh, Amar Phanishayee, Onur Mutlu, and Phillip B. Gibbons. The non-iid data quagmire of decentralized machine learning, 2020.
- [17] Qinbin Li, Yiqun Diao, Quan Chen, and Bingsheng He. Federated learning on non-iid data silos: An experimental study. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*, pages 965–978, 2022.
- [18] Sashank Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný, Sanjiv Kumar, and H. Brendan McMahan. Adaptive federated optimization, 2021.
- [19] Xiaodong Ma, Jia Zhu, Zhihao Lin, Shanxuan Chen, and Yangjie Qin. A state-of-the-art survey on solving non-iid data in federated learning. *Future Generation Computer Systems*, 135:244–258, 2022.
- [20] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- [21] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986.
- [22] Benjamin Schumacher and Michael Westmoreland. *Quantum processes systems, and information*. Cambridge University Press, 2010.
- [23] W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, 1982.
- [24] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii. Quantum circuit learning. *Physical Review A*, 98(3), September 2018.
- [25] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. Evaluating analytic gradients on quantum hardware. *Physical Review A*, 99(3), March 2019.
- [26] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. Evaluating analytic gradients on quantum hardware. *Physical Review A*, 99(3), March 2019.
- [27] James C. Spall. An overview of the simultaneous perturbation method for efficient optimization. *Johns Hopkins Apl Technical Digest*, 19:482–492, 1998.
- [28] James C. Spall. Implementation of the simultaneous perturbation algorithm for stochastic optimization. *IEEE Transactions on Aerospace and Electronic Systems*, 34:817–823, 1998.
- [29] Gokul Subramanian Ravi, Tan Nguyen, Robert Rines, Gideon Chern, Thomas O’Brien, and Jay M. Gambetta. Quantum computing in the cloud: Analyzing job and machine characteristics. *arXiv preprint arXiv:2203.13121*, 2022. Analyzes large-scale IBM Quantum workloads and execution time metrics in cloud settings.
- [30] Darshana Lall, Stephen Kelly, Neil Bingham, Alexander Cowtan, Jaspreet Singh, and Ross Duncan. A review and collection of metrics and benchmarks for quantum computers: Definitions, methodologies, and software. Technical report, University of Strathclyde, 2025. Provides standardized runtime and performance metrics for quantum processors.
- [31] Ligeng Zhu, Zhijian Liu, and Song Han. Deep leakage from gradients. *Advances in neural information processing systems*, 32, 2019.
- [32] Ali Javadi-Abhari, Matthew Treinish, Kevin Krsulich, Christopher J. Wood, Jake Lishman, Julien Gacon, Simon Martiel, Paul D. Nation, Lev S. Bishop, Andrew W. Cross, Blake R. Johnson, and Jay M. Gambetta. Quantum computing with Qiskit, 2024.
- [33] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

- [34] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- [35] Sukin Sim, Peter D. Johnson, and Alán Aspuru-Guzik. Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. *Advanced Quantum Technologies*, 2(12), October 2019.
- [36] Maria Schuld, Ryan Sweke, and Johannes Jakob Meyer. Effect of data encoding on the expressive power of variational quantum-machine-learning models. *Physical Review A*, 103(3), March 2021.